



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
-----------------	-------------	----------------------	---------------------	------------------

10/608,139

06/30/2003

Sanjay Ghemawat

0026-0029

3005

44989 7590 08/26/2009

HARRITY & HARRITY, LLP
11350 Random Hills Road
SUITE 600
FAIRFAX, VA 22030

EXAMINER

THAI, HANH B

ART UNIT

PAPER NUMBER

2163

MAIL DATE

DELIVERY MODE

08/26/2009

PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE BOARD OF PATENT APPEALS
AND INTERFERENCES

Ex parte SANJAY GHEMAWAT,
HOWARD GOBIOFF,
and SHUN-TAK LEUNG

Appeal 2008-004150
Application 10/608,139¹
Technology Center 2100

Decided: August 26, 2009

Before LEE E. BARRETT, LANCE LEONARD BARRY, and
THU A. DANG, *Administrative Patent Judges*.

BARRETT, *Administrative Patent Judge*.

DECISION ON APPEAL

¹ Filed June 30, 2003, titled "Systems and Methods for Replicating Data," which claims the benefit of Provisional Applications 60/447,277, filed February 14, 2003, and 60/459,648, filed April 3, 2003.

This is a decision on appeal under 35 U.S.C. § 134(a) from the final rejection of claims 1-8 and 19-27. Claims 9-18 stand withdrawn as a result of an election to a restriction requirement. We have jurisdiction pursuant to 35 U.S.C. § 6(b).

We reverse.

STATEMENT OF THE CASE

The invention

The invention relates to the distribution and redistribution of chunks of data among multiple servers. A chunk is a fixed-size portion of a file. For reliability, each chunk may be replicated on multiple chunk servers. The invention relates to how to select servers to store a replica of the data and how to redistribute replicas among the servers. *See* Abstract.

The representative claims

Claims 1 and 19 are reproduced below:

1. A method for distributing data in a system that includes a plurality of servers, the method comprising:

identifying ones of the servers to store a replica of the data based on at least one of utilization of the servers, prior data distribution involving the servers, or failure correlation properties associated with the servers; and

placing the replicas of the data at the identified servers.

19. A method for redistributing chunks of data in a system that includes a plurality of servers that store replicas of the chunks, the method comprising:

monitoring utilization of the servers;

determining whether to redistribute any of the replicas;

selecting one or more of the replicas to redistribute based on the utilization of the servers;

selecting one or more of the servers to which to move the one or more replicas; and

moving the one or more replicas to the selected one or more servers.

The references

Narendran	US 6,070,191	May 30, 2000
Jindal	US 6,324,580 B1	Nov. 27, 2001

The rejections²

Claims 1-8 stand rejected under 35 U.S.C. § 102(b) as being anticipated by Jindal.

Claims 19-27 stand rejected under 35 U.S.C. § 103(a) as unpatentable over Jindal and Narendran.

DISCUSSION

Anticipation - claims 1-8

Issues

Based on Appellants' arguments, the issues are: Have Appellants shown that the Examiner erred in finding that Jindal teaches the limitations of: (1) "identifying ones of the servers to store a replica of the data based on at least one of utilization of the servers, prior data distribution involving the servers, or failure correlation properties associated with the servers"; and (2) "placing the replicas of the data at the identified servers," as recited in claim 1? Claim 7 contains corresponding limitations in means-plus-function

² The rejection of claims 1-8 and 19-27 under 35 U.S.C. § 101 is withdrawn in the Examiner's Answer (Ans. 9).

language, and claim 8 contains corresponding limitations in a file system; thus, claim 1 is representative of claims 1, 7, and 8.

Principle of law

"Anticipation requires the presence in a single prior art disclosure of all elements of a claimed invention arranged as in the claim." *Connell v. Sears, Roebuck & Co.*, 722 F.2d 1542, 1548 (Fed. Cir. 1983).

Contentions

The Examiner finds that Jindal teaches the "identifying" limitation at the Abstract, Summary, column 4, lines 40-67, and column 5, lines 57-60 and teaches the "placing" limitation at the Abstract, Summary, and column 4, lines 40-67. Final Office Action 4.

Appellants discuss each of these portions of Jindal in detail and argue that none discloses or remotely suggests "identifying ones of the servers to store a replica of the data based on at least one of utilization of the servers, prior data distribution involving the servers, or failure correlation properties associated with the servers" as recited in claim 1. Br. 32-35.

Appellants further argue that the Examiner-cited portions of Jindal do not disclose or remotely suggest "placing the replicas of the data at the identified servers" as recited in claim 1. Br. 35-36.

The Examiner responds:

Jindal clearly discloses that executing a replicated service is selected and processed to a preferred server (col. 4, lines 49-67; col. 5). This teaching of Jindal is clearly illustrating the storing of replicas of the data as demonstrated in applicant's claimed language of identifying and placing the replicas of the data at the identified server.

Ans. 10.

Appellants argue that Jindal provides absolutely no support for the Examiner's allegation. It is noted that column 4 of Jindal discloses that a replicated service is available on multiple servers, each of which operates a separate instance of the replicated service, and discloses identifying a preferred server, as a least-loaded or closest server, to which to direct subsequent requests for the replicated service. Reply Br. 4. It is argued that this is completely different from identifying servers to store a replica of data, let alone identifying servers to store a replica based on one of the claimed criteria. *Id.* at 4-5. Appellants argue that column 5 of Jindal discloses balancing requests for a replicated service among multiple servers, which is a completely different principle from identifying servers to store a replica of data, let alone identifying servers to store a replica based on one of the claimed criteria. *Id.* at 5.

Appellants argue that the Examiner fails to address the arguments about placing the replica of the data at the identified servers, and thus apparently concedes the validity of the arguments. Reply Br. 5.

Facts

Jindal describes in the Background that a service (or application) offered simultaneously on multiple servers is termed "replicated" in recognition of the fact that each instance of the service operates in substantially the same manner and provides substantially the same functionality as the others. Col. 1, ll. 18-22. The method of distributing client requests for a service or application among the servers is known as "load balancing." Col. 1, ll. 22-28.

Jindal discloses a method for load balancing requests for a replicated service or application among a plurality of servers operating instances of the replicated service or application. Abstract.

Jindal describes that a policy is selected for identifying a "preferred server" according to status or operational characteristics, such as the least-loaded or closest server. Abstract; col. 4, ll. 40-67; col. 5, ll. 57-60. Information used to identify the preferred server is gathered by status objects. Col. 6, ll. 32-46 and 56-64. Requests for the replicated service or application are then directed to the preferred server until a different preferred server is identified. Abstract.

Analysis

Appellants' arguments are persuasive. Jindal describes identifying a preferred server to handle clients' requests for a replicated service or application, not identifying a server "to store a replica of the data" as claimed. The Examiner points to nothing in Jindal that describes storing data, much less a replica of data, and we find no teaching of storing data. To the extent the Examiner may be interpreting the replicated service or application as a "replica of the data," this interpretation is not expressed or implied in the rejection, but in any case, Jindal does not describe how the servers were selected to store a replicated service or application. We agree with Appellants that identifying a preferred server to which to direct a client request for a replicated service or application is a completely different function from identifying a server to store a replica of data. While the step of identifying a preferred server based on the "policy" of the "least-loaded" server in Jindal (col. 4, ll. 60-61) sounds like identifying a server "based on

utilization of the servers," as recited in claim 1, Jindal does not describe the function of identifying a server to "store a replica of data."

The portions of Jindal relied upon by the Examiner do not support the Examiner's position for the reasons stated by Appellants. The Abstract, Summary, column 4, and column 5, all teach identifying a preferred server to handle requests for a replicated service or application, not identifying a server "to store a replica of the data."

Appellants have shown that the Examiner erred in finding that Jindal describes "identifying ones of the servers to store a replica of the data based on at least one of utilization of the servers, prior data distribution involving the servers, or failure correlation properties associated with the servers."

As to the limitation of "placing the replicas of the data at the identified servers," none of the portions of Jindal cited by the Examiner describes or suggests this limitation. The Examiner's rejection does not explain how the cited portions of Jindal in any way teach the limitation. Jindal describes identifying and then directing requests to a preferred server, not putting replicas of data on any of the servers.

Thus, Appellants have also shown that the Examiner erred in finding that Jindal describes "placing the replicas of the data at the identified servers."

Conclusion

Appellants have shown that the Examiner erred in finding that Jindal teaches the limitations of: (1) "identifying ones of the servers to store a replica of the data based on at least one of utilization of the servers, prior data distribution involving the servers, or failure correlation properties

associated with the servers"; and (2) "placing the replicas of the data at the identified servers," as recited in representative claim 1. The rejection of claims 1-6 is reversed. Claims 7 and 8 contain corresponding limitations. Accordingly, the rejection of claims 7 and 8 is reversed.

Obviousness - claims 19-27

Issue

Based on Appellants' arguments, the issues are: Have Appellants shown that the Examiner erred in concluding that the combination of Jindal and Narendran disclose or suggest the limitations of: (1) "determining whether to redistribute any of the replicas"; (2) "selecting one or more of the replicas to redistribute based on the utilization of the servers"; (3) "selecting one or more of the servers to which to move the one or more replicas"; and (4) "moving the one or more replicas to the selected one or more servers" as further recited in claim 19? Claim 26 contains corresponding limitations in means-plus-function language and claim 27 contains corresponding limitations as functions of a file system; thus, claim 19 is representative of independent claims 19, 26, and 27.

Principles of law

For an obviousness rejection, the combination of references must teach or suggest to a person of ordinary skill in the art all of the claim limitations. 35 U.S.C. § 103(a).

Contentions

The Examiner finds that Jindal discloses "determining whether to redistribute any of the replicas" and "selecting one or more of the replicas to

redistribute based on the utilization of the servers" at column 6, lines 31-45 and 56-64. Final Office Action 6. The Examiner finds that Jindal discloses the limitations of claim 19 except that it "does not explicitly disclose redistribute the replicas." *Id.* The Examiner finds that "Narendran . . . discloses data distribution techniques for load-balanced fault-tolerant web access including redistributing replicated data from the failed server to achieve rebalance (see col. 12, lines 12-38)." *Id.* at 6-7. The Examiner concludes that it would have been obvious "to utilize the redistributing or rebalance technique of Narendran to derive the invention as claimed. The motivation of doing so would have been to provide an efficient system that can maintain a maximum flow with a minimum cost of a network flow system (abstract of Narendran)." *Id.* at 7.

Appellants argue that Jindal does not disclose "determining whether to redistribute any of the replicas." It is argued that column 6, lines 31-45, discloses choosing a least-loaded server or a closest server, not whether to redistribute any of the replicas. Br. 45-46. It is argued that column 6, lines 56-64, discloses a status object for each replicated service or application to be monitored, and does not suggest determining whether to redistribute any of the replicas. *Id.* at 46. Appellants argue that the Examiner admitted that Jindal does not disclose redistributing replicas and therefore no reasonable argument can be made that it discloses determining whether to distribute any of the replicas. *Id.*

The Examiner responds:

The replicated system [in Jindal] has to have the ability of redistribution despite that it does not recite the claimed limitation in verbatim. Jindal clearly discloses the replicated services include the

redistribution based on utilization of server (i.e. "referred [sic, preferred] server") as illustrated in the claimed language.

Ans. 12. The Examiner also reiterates the finding and conclusion regarding Narendran. *Id.* at 12-13.

Appellants argue that the Examiner's statements are contrary to prior admissions. Reply Br. 11. Appellants argue that Jindal discloses identifying a server to which to direct a request for a replicated service, which is completely different from determining whether to redistribute any of the replicas of chunks of data. *Id.* at 11-12.

Appellants argue that Narendran does not describe redistributing replicated data as stated by the Examiner, but "discloses rebalancing server loads by adjusting the redirection probabilities used by the redirection server when a server fails, changes in the access probabilities of the document occurs, or changes in the capacity of the document servers occurs" (Br. 47), i.e., Narendran teaches away because it states that rebalancing should be accomplished without redistributing the documents. It is argued that Narendran teaches away as previously argued, and that the Examiner's statement of motivation does not explain how redirection of data would improve a computer's efficiency. Reply Br. 12-13.

Appellants further argue that Jindal and Narendran also do not disclose "selecting one or more of the replicas to redistribute based on the utilization of the servers." It is argued that Jindal does not disclose this limitation at column 6 for reasons similar to those discussed. Br. 49.

Appellants argue that because Jindal and Narendran do not disclose or suggest "determining whether to redistribute any of the replicas" or "selecting one or more of the replicas to redistribute based on the utilization

of the servers," the combination cannot reasonably disclose or suggest "selecting one or more of the servers to which to move the one or more replicas" or "moving the one or more replicas to the selected one or more servers" as further recited in claim 19. Br. 49.

Facts

Narendran is directed to a server system for processing client requests for documents distributed across a document server cluster. Abstract.

Narendran describes that a redirection server receives a client request and redirects it to one of the document servers based on a set of pre-computed redirection probabilities. Abstract.

Narendran describes that a load distribution algorithm attempts to equalize the sum of the access rates of all the documents stored at a given document server across all of the document servers. Abstract.

Narendran describes an algorithm, referred to as a "binning" algorithm, used to provide an initial document distribution that achieves load balance in a heterogeneous server cluster. Col. 5, l. 47 to col. 6, l. 27.

Narendran describes that if the load balance achieved through initial distribution of the document is disturbed, it is desirable to rebalance the server loads "without redistributing documents among the servers, since this would generally involve overhead in moving the documents between servers and may affect availability of the system. It is preferable to instead achieve rebalance by adjusting only the redirection probabilities used by the redirection server." Col. 12, ll. 23-28.

Analysis

Appellants' arguments are persuasive. Jindal describes identifying a preferred server to handle clients' requests for a replicated service or application. Thus, Jindal may disclose determining whether to redistribute clients' *requests* for a replication service or application, not "determining whether to redistribute any of the replicas [of chunks]." Requests for a service are not replicas of chunks, as claimed. If the Examiner considers the replicated services and applications in Jindal to be the claimed "replicas," because of the similarity in terminology, we note that Jindal does not describe how the replicated service or application are initially distributed, and certainly does not describe determining whether to redistribute them.

Appellants have shown that the Examiner erred in finding that Jindal describes "determining whether to redistribute any of the replicas."

As to the limitation of "selecting one or more of the replicas to redistribute based on the utilization of the servers," Jindal does not describe or suggest this limitation, which is not surprising since it does not redistribute replicas. Narendran does not describe redistributing replicated data. Narendran describes that if the load balance of client requests for documents achieved through initial distribution of the document is disturbed, it is desirable to rebalance the server loads "without redistributing documents among the servers, since this would generally involve overhead in moving the documents between servers and may affect availability of the system. It is preferable to instead achieve rebalance by adjusting only the redirection probabilities used by the redirection server." Col. 12, ll. 23-28. While this does not teach away from moving documents, but only that moving documents is not preferred, Narendran does not teach "selecting one or more of the replicas to redistribute based on the utilization of the servers."

Thus, Appellants have also shown that the Examiner erred in finding that the combination of Jindal and Narendran describes "selecting one or more of the replicas to redistribute based on the utilization of the servers."

Manifestly, since Jindal and Narendran do not disclose or suggest "determining whether to redistribute any of the replicas" or "selecting one or more of the replicas to redistribute based on the utilization of the servers," the combination cannot reasonably disclose or suggest "selecting one or more of the servers to which to move the one or more replicas" or "moving the one or more replicas to the selected one or more servers" as further recited in claim 19.

Conclusion

Appellants have shown that the Examiner erred in concluding that the combination of Jindal and Narendran discloses or suggests the limitations of: (1) "determining whether to redistribute any of the replicas"; and (2) "selecting one or more of the replicas to redistribute based on the utilization of the servers"; (3) "selecting one or more of the servers to which to move the one or more replicas"; and (4) "moving the one or more replicas to the selected one or more servers," as recited in representative claim 19. The rejection of claims 19-25 is reversed. Independent claims 26 and 27 contain corresponding limitations. Accordingly, the rejection of claims 26 and 27 is also reversed.

Appeal 2008-004150
Application 10/608,139

CONCLUSION

The rejections of claims 1-8 and 19-27 are reversed.

REVERSED

erc

HARRITY & HARRITY, LLP
11350 Random Hills Road
SUITE 600
FAIRFAX, VA 22030